Federal Agencies Prospective on High End Computing

David B. Nelson, Ph.D.
Director
National Coordination Office for
Information Technology Research and Development

April 8, 2003

IDC HPC User Forum Meeting

Networking and Information Technology Research and Development (NITRD) Program

- Coordinated, focused long-term interagency R&D in information technologies
- Evolved from the Federal HPCC, CIC, NGI, and IT R&D programs
- $2 billion multi-agency NITRD Program
  - 12 agencies and departments coordinated via a "virtual agency" coordination/management structure
    - AHRQ, DARPA, DOE/NNSA, DOE/SC, EPA, NASA, NIH, NIST, NOAA, NSA, NSF, ODDR&E
  - Coordinated by the National Coordination Office for Information Technology Research and Development
- Assessed by the President's Information Technology Advisory Committee

NITRD Program Coordination

WHITE HOUSE

Executive Office of the President
Office of Science and Technology Policy

National Science and Technology Council

President's Information Technology Advisory Committee (PITAC)

U.S. Congress

NITRD Authorization and Appropriations Legislation

National Coordination Office (NCO) for Information Technology Research and Development

Interagency Working Group (IWG) on Information Technology R&D

Participating Agencies: AHRQ, DARPA, DOE/NNSA, DOE/SC, EPA, NASA, NIH, NIST, NOAA, NSA, NSF, ODDR&E

- **High End Computing Coordinating Group (HEC)**
- **Large Scale Networking Coordinating Group (LSN)**
- **High Confidence Software and Systems Coordinating Group (HCSS)**
- **Human Computer Interaction & Information Management Coordinating Group (HCI & IM)**
- **Software Design and Productivity Coordinating Group (SDP)**
- **Social, Economic and Workforce Implications of IT and IT Workforce Development Coordinating Group (SEW)**

National Coordination Office (NCO) for Information Technology Research and Development (IT R&D)

- NCO Director reports to the Director of the White House Office of Science Technology Policy (OSTP)
- Coordinates planning, budget, and assessment activities for the Federal multiagency Networking and Information Technology Research and Development (NITRD) Program
- Supports the six technical Coordinating Groups (CGs) that report to the Interagency Working Group (IWG) on IT R&D
  - Research planning workshops, conferences, and meetings
  - Presentations, white papers, and research reports
- Provides technical and administrative support to the IWG and to the President's Information Technology Advisory Committee (PITAC)
- Informs the public of Federal achievements and challenges in IT R&D
  - Maintains a Web site
  - Publishes annual budget documents in cooperation with the IT R&D agencies
  - Publishes PITAC reports

Performance Measures of Selected Top Computers

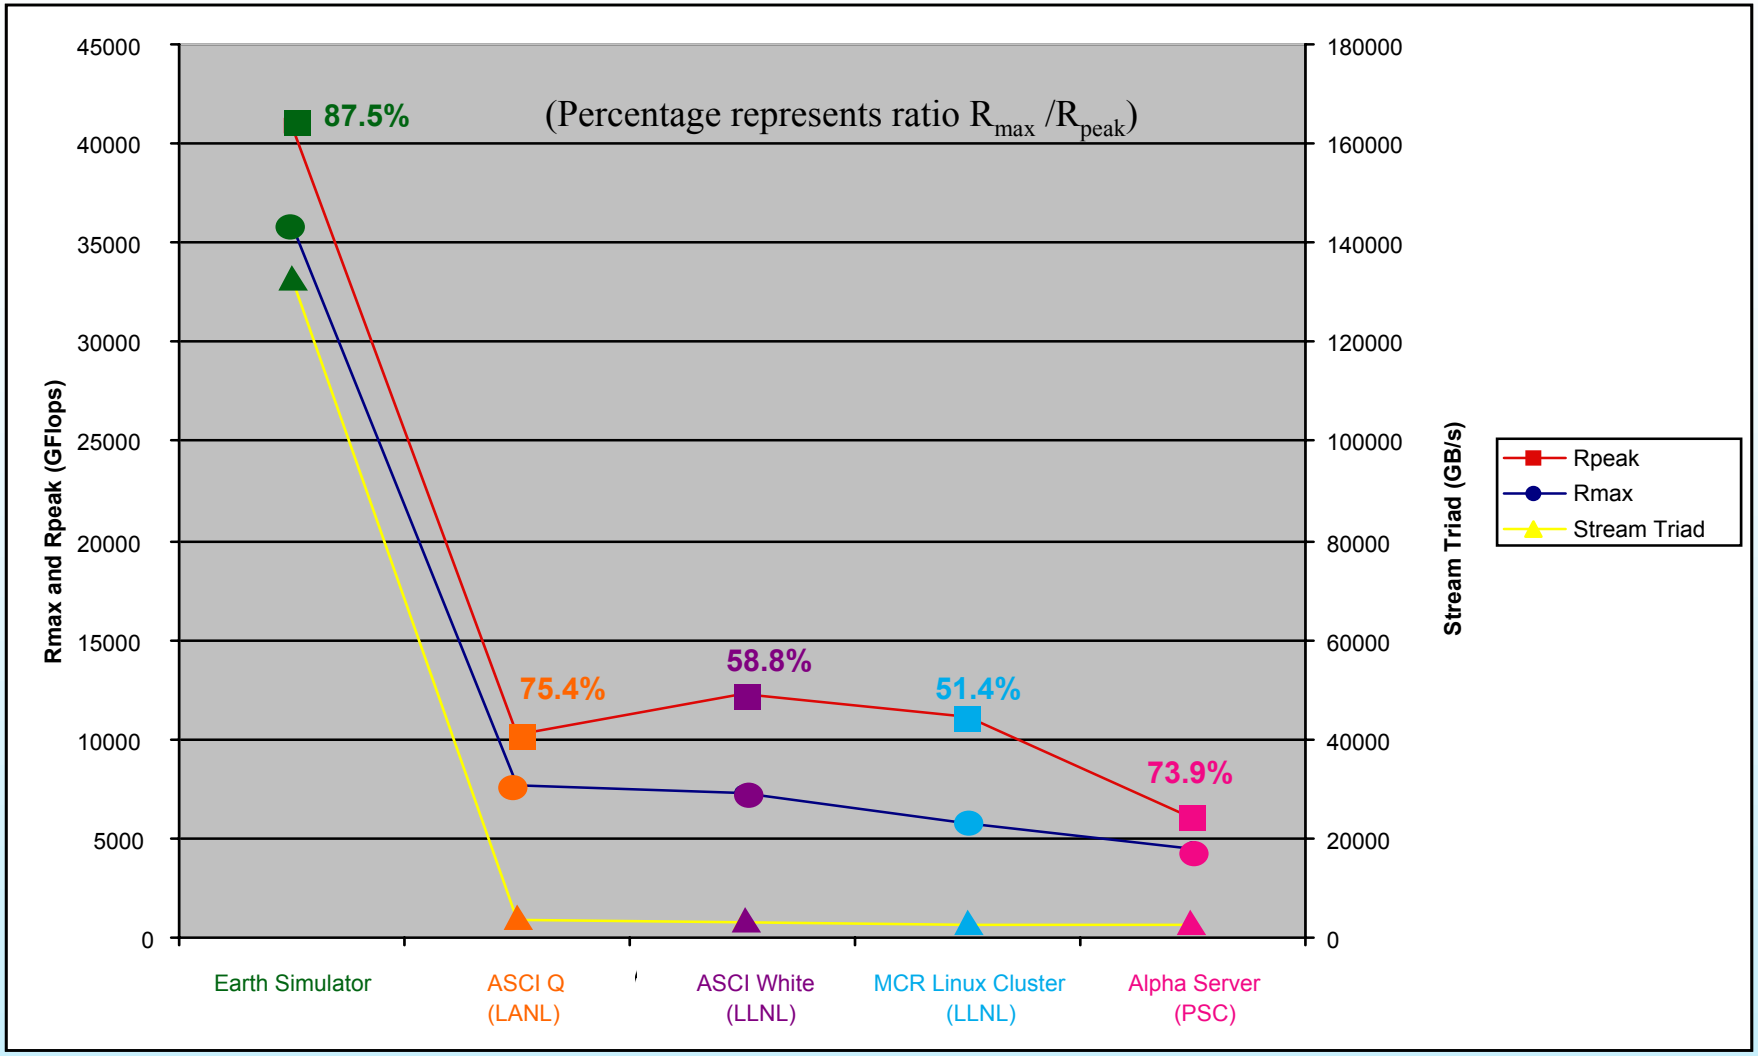

R_{peak} and R_{max} data from Top500.org Stream Triad data from IDC

Several Federal Agencies Have Recently Examined High End Computing Needs

- Agencies include:
  - DARPA, DOE/NNSA, DOE/SC, EPA, NASA, NIH, NIST, NOAA, NSA, NSF, ODDR&E
- They are mostly using COTS-based HEC
- Most expect COTS to be acceptable in near term, however:
  - Time-to-solution becoming too long
    - Too hard to program; too hard to optimize
    - Need improved programming environments
  - For many problems, computing efficiency is dropping as number of processors increases
    - Affects time-to-solution, cost-to-solution
  - Memory performance (latency/bandwidth/size) is often the limiting factor
  - For some applications I/O system performance is limiting

Examples of HEC Application Area With Federal Interest

- Nuclear stockpile stewardship (multi-discipline physics models)
- Global and regional climate modeling
- Weather and Ocean Forecasting
- Geophysics (earthquakes, volcanoes, landslides, plate tectonics, magneto-dynamics)
- Astrophysics (star and galaxy dynamics)
- Aeronautics and aerospace design (air-frames, re-entry vehicles)
- Engineering design of ships, land vehicles, buildings
- Weapon designs and weapon effects
- Armor design
- Survivability/stealthiness design
- Signal and imagining processing
- Signals intelligence (e.g. cryptanalysis)
- Electromagnetics
- Molecular modeling for chemical risk assessment
- Biophysics (e.g., protein folding)
- Pharmacology
- Quantum chemistry
- Materials modeling and design (e.g., concrete)
- Quantum chromodynamics

For Some Important Applications/Algorithms COTS-based HEC May be Unsuitable

- Primarily due to non-local memory reference (e.g., long vectors requiring gather-scatter operations)

- Fixes would require some mix of :
  - Better memory subsystems performance (latency, bandwidth, size)
  - Better ability to handle large vector operations including gather-scatter

Examples of Applications for Which COTS May be Unsuitable

- Hypersonic air-breathing propulsion
  - Needs high memory-to-CPU bandwidth for multi-disciplinary analysis
- Reusable launch vehicle design
  - Needs high memory-to-CPU bandwidth
- Protein folding
  - Some algorithms are poorly parallelizable
- Cryptoanalysis
  - Needs fast flat-memory model
- Climate data assimilation
  - Part of problem not easily parallelizable, needs high memory-to-CPU bandwidth

Agency Conclusions

- Further progress in HEC will require balanced, coordinated effort in:
  - Research, development, and engineering of new HEC architectures and systems
  - Procurement of new COTS and custom systems
  - Better software (systems, middleware, and applications)
  - Better domain science (mathematics and algorithms)
- HEC is a decreasing part of the technical computing marketplace.
- COTS-based HEC is largely based on technologies developed for low- and mid-range markets (SMP nodes, low bandwidth interconnects).
- Market pressure may result in future COTS systems being less responsive to HEC needs.
- Federal funding of highest-performing HEC, including development of new systems, may be required.

High End Computing (HEC) in President's FY 2004 Budget

"Due to its impact on a wide range of federal agency missions ranging from national security and defense to basic science, high end computing—or supercomputing—capability is becoming increasingly critical. Through the course of 2003, agencies involved in developing or using high end computing will be engaged in planning activities to guide future investments in this area, coordinated through the NSTC. The activities will include the development of interagency R&D roadmap for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues (along with recommendations where applicable) relating to federal procurement of high-end computing systems. The knowledge gained for this process will be used to guide future investments in this area. Research and software to support high end computing will provide a foundation for future federal R&D by improving the effectiveness of core technologies on which next-generation high-end computing systems will rely."

High End Computing Revitalization Task Force (HECRTF)

- Rationale: High End Computing (HEC) increasingly critical
- HECRTF commissioned by Office of Science and Technology Policy and coordinated through National Science and Technology Council (NSTC)
- HECRTF will develop a five-year plan to guide future Federal HEC investments
- Plan will lay out an overall strategy for these investments
- Wide participation in HECRTF by Federal agencies developing or using HEC
- Final report to be completed by August 2003, in time to serve as input to FY 2005 budget
  - May 1: Rough draft
  - June 1: Draft
  - August: Final report
- Outreach strategy being developed,will likely include
  - Request for white papers
  - Meetings
  - Major workshop

HECRTF Organization

- Task Force - Oversight and guidance
- Task Groups - Specific issue
  - Integration
  - Core Technologies Research & Development
  - Capacity, Capability, and Accessibility
  - Procurement of Federal HEC Systems
- Coordination and support from NCO

Report of the Panel on Large Scale Computing in Science and Engineering (Lax Panel) Dec. 26, 1982

"Those who do not remember the past are condemned to repeat it"
George Santayana*

Recommendations:

- The Panel recommends that a long-term National Program on Large Scale Computing should be initiated immediately, with the participation of the appropriate Federal agencies, the universities, and industry.

The goals of this National Program should be:

- Increased **access** for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputing facilities and experimental computers;
- Increased **research** in computational mathematics, software , and algorithms necessary to the effective and efficient use of supercomputer systems;
- **Training** of personnel in scientific and engineering computing and
- Research and development basic to the **design** and implementation of new supercomputer systems of substantially increased capability beyond that likely to arise from commercial sources.

This program should be coordinated by an interagency policy committe consisting of representative of the appropriate Federal agencies, including DOC, DOD, DOE, NASA, and NSF. A Large Scale Computing Advisory Panel, with representatives from Government, the universities, and industry, should be established to assist in the planning, implementation and operation of the Program.

Thanks to Norm Kreisman, DOE, for inspiration and reference material

A Research and Development Strategy for High Performance Computing, Nov. 20, 1987

Summary of Recommendation for a National High Performance Computing Strategy:

1 **HIGH PERFORMANCE COMPUTERS:** The U.S. Government should establish a long range strategy for Federal support for basic research on high performance computer technology and the appropriate transfer of research and technology to U.S. industry.

2 **SOFTWARE TECHNOLOGY AND ALGORITHMS:** The U.S. should take the lead in encouraging joint research with government, industry, and university participation to improve basic tools, languages, algorithms, and associated theory for the scientific "grand challenges" with widespread applicability.

3 **NETWORKING:** U.S. government, industry, and universities should coordinate research and development for a research network to provide a distributed computing capability that links the government, industry, and higher education communities.

4 **BASIC RESEARCH AND HUMAN RESOURCES:** Long term support for basic research in computer science should be increased within available resources. Industry, universities, and government should work together to improve the training and utilization of personnel to expand the base of research and development in computational science and technology .

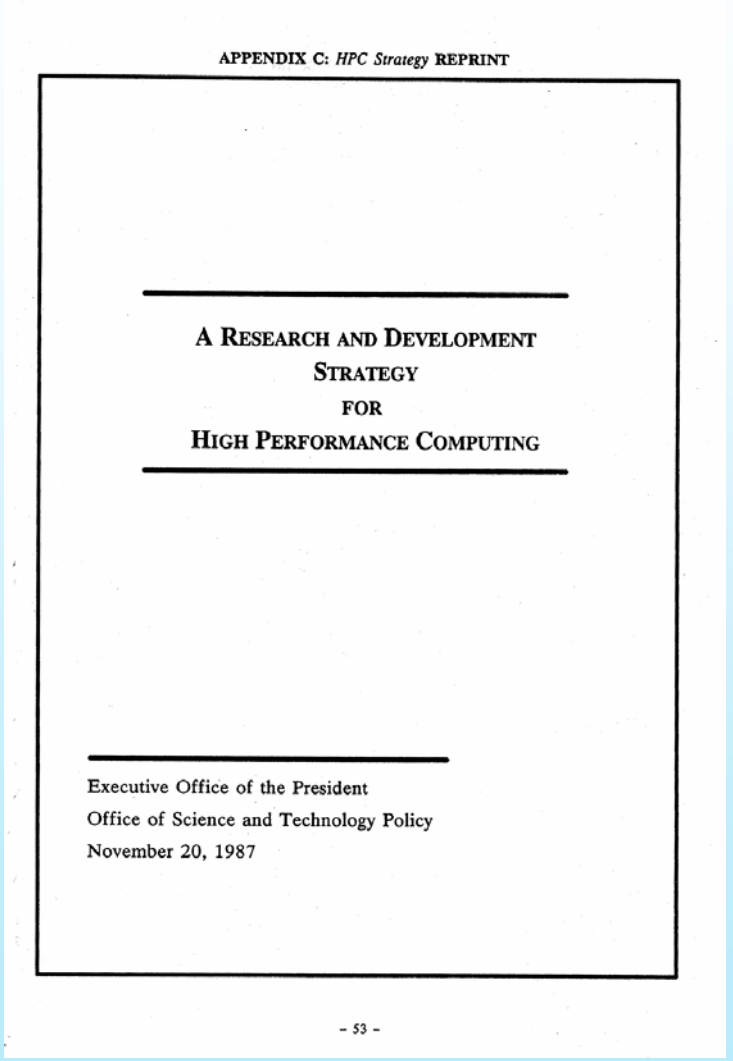

APPENDIX C: *HPC Strategy* REPRINT

A RESEARCH AND DEVELOPMENT STRATEGY FOR HIGH PERFORMANCE COMPUTING

Executive Office of the President
Office of Science and Technology Policy
November 20, 1987

- 53 -

The U.S. Supercomputer Industry, Federal Coordinating Council or Science Engineering and Technology (FCCSET) December 1987

FCCSET Recommends that:

1. The U.S. Government carry out a coordinated long-range R&D program in supercomputer applications, software, and advanced computer architectures

2. The Federal Government should return its role as a “friendly” buyer of innovative supercomputers

3. The Federal Government should make federally developed software available to the private sector when possible

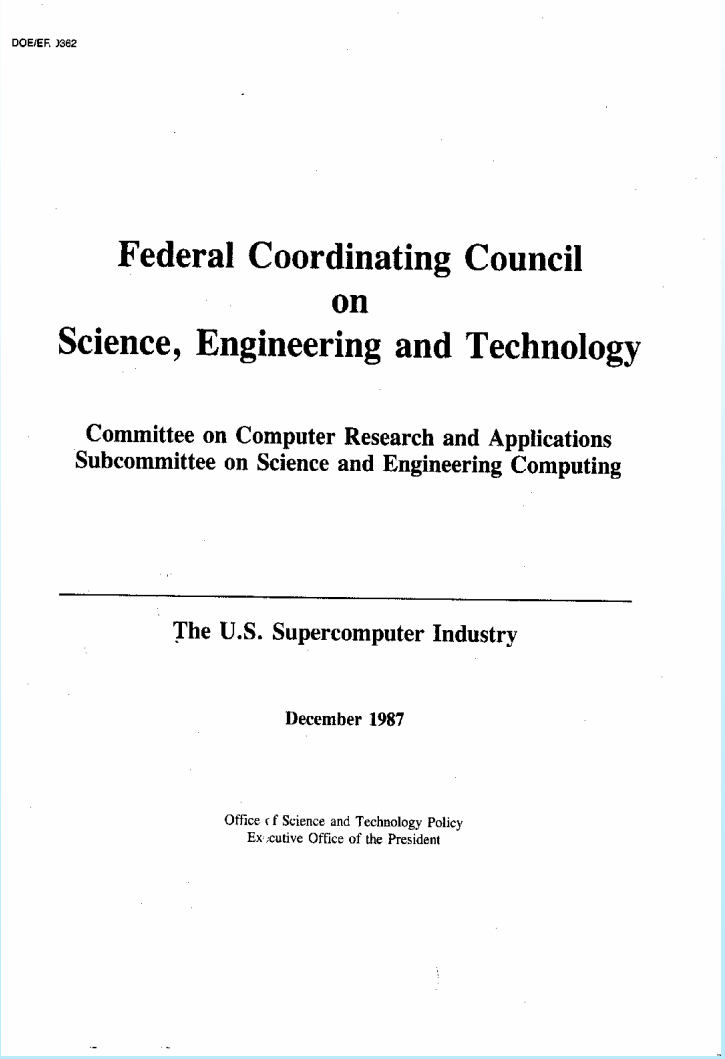

DOE/ER 0362

Federal Coordinating Council
on
Science, Engineering and Technology

Committee on Computer Research and Applications
Subcommittee on Science and Engineering Computing

The U.S. Supercomputer Industry

December 1987

Office of Science and Technology Policy
Executive Office of the President

The Federal High Performance Computing Program, OSTP, September 8, 1989

GOALS:

- **Maintain and extend U.S. leadership in high performance computing, and encourage U.S. sources of production;**

- **Encourage innovation in high performance computing technologies by increasing their diffusion and assimilation into the U.S. science and engineering communities; and**

- **Support U.S. economic competitiveness and productivity through greater utilization of networked high performance computing in analysis, design, and manufacturing**

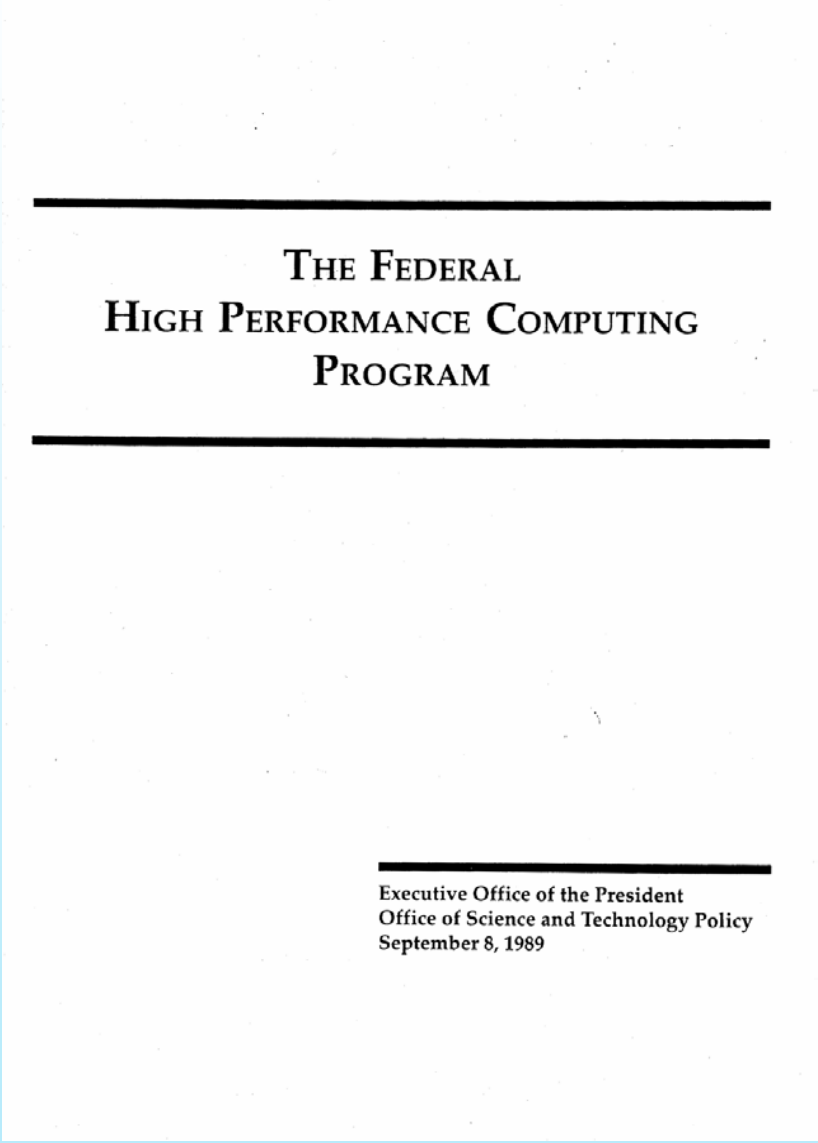

THE FEDERAL
HIGH PERFORMANCE COMPUTING
PROGRAM

Executive Office of the President
Office of Science and Technology Policy
September 8, 1989

For Further Information

Please contact us at:

nco@itrd.gov

Or visit us on the Web:

www.itrd.gov